

U.S. PATENT APPLICATION

Inventor(s): Matthew T. HART

Invention: UNWANTED E-MAIL FILTERING

*NIXON & VANDERHYE P.C.
ATTORNEYS AT LAW
1100 NORTH GLEBE ROAD
8TH FLOOR
ARLINGTON, VIRGINIA 22201-4714
(703) 816-4000
Facsimile (703) 816-4100*

SPECIFICATION

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION PAPERS

OF

MATTHEW THOMAS HART

FOR

UNWANTED E-MAIL FILTERING

BACKGROUND OF THE INVENTION

Field of the Invention

This invention relates to the field of data processing systems. More particularly, this invention relates to the field of e-mail filtering within such data processing systems.

With the rise in the use of e-mail as a communication mechanism, this has been accompanied by a rise in the occurrence of unsolicited and unwanted e-mail messages. These so-called "Spam" messages cause a number of problems, such as consuming physical network and processing resources as well as wasting the time of the recipients in dealing with these messages in their inbox.

Description of the Prior Art

It is known to provide e-mail filtering mechanisms that apply predefined rules to received e-mail messages in order that Spam messages may be identified and automatically deleted. These existing system often work on content filtering with rules based on regular expressions applied to all inbound messages. A significant amount of unwanted e-mail is generally still able to pass through such systems because the filtering rules are not normally maintained to a high degree and tend to lack flexibility to deal with an ever changing problem.

SUMMARY OF THE INVENTION

Viewed from one aspect the present invention provides a computer program product comprising a computer program operable to control a computer to process received e-mail messages, said computer program comprising:

- (i) filter downloading logic operable to download filter data from a remote source, said filter data specifying a plurality of tests that may be used to identify unwanted e-mail messages;
- (ii) e-mail filtering logic operable to receive an e-mail message and to apply said plurality of tests to identify unwanted e-mail messages; and
- (iii) unwanted message reporting logic operable to allow reporting to a filter data generator a new unwanted e-mail message received and not identified by

said plurality of tests such that said filter data may be updated to identify said new unwanted e-mail message.

The invention recognises that unwanted e-mail messages are not generally restricted to a single user and that filtering rules developed in response to receipt of an unwanted e-mail message by one user may well be of use to another user who has yet to receive any of that unwanted e-mail. The invention also recognises that the value of allowing users to report the receipt of new unwanted e-mail messages not already trapped by the filters is that the positive identification of that mail as wanted by a user is very strongly indicative of the mail genuinely being a Spam mail that will be unwanted by all users. This contrasts with computer virus reporting or bug reporting by users where the updating of a central resource by a provider needs much more careful consideration before being performed as users may often be incorrect in their assessment of the nature of the problem. Compared to this, whether or not an e-mail is an unwanted e-mail is a decision that is primarily made in the mind of the recipient and so a report of such an e-mail message to a provider of filtered data is substantially definitive in establishing that the filter data should be modified to prevent receipt of that unwanted e-mail message. This lends the mechanisms well suited to being substantially automated thereby giving a possibility of faster filter generation and anti-Spam protection.

The tests applied by the filtering mechanisms of preferred embodiments use scoring algorithms to identify received e-mail messages as unwanted e-mail messages. The scoring algorithms are generally more flexible and have a chance of identifying new unwanted e-mail messages at their first occurrence due to content matching known criteria for unwanted e-mail messages, such as the presence of predetermined words in a high proportion or characteristics of the addressee list.

A particularly preferred feature of the invention is that should the tests identify an e-mail message as potentially unwanted then it is forwarded to its addressee together with a prompt that allows the addressee to provide feedback as to whether or not in their opinion the e-mail is an unwanted e-mail message. This preferred feature builds upon the realisation that the determination of whether or not an e-mail message is an unwanted e-mail message is primarily in the mind of the recipient and

accordingly allowing the recipient to make this decision enables the load of maintaining the rules set to be distributed and a faster and more reliable response achieved.

A particularly preferred way of facilitating such feedback is to encapsulate the suspect e-mail message within a mark up language document that provides voting buttons to allow the addressee to give their feedback to the system.

Whilst the system could be arranged such that new rules could only be created centrally within the downloading source, preferred embodiments provide the ability for local rules to be created. This allows a faster response for an organisation receiving problems through unwanted e-mail messages and also allows an organisation to treat as unwanted e-mail messages that may not qualify as such in the view of the provider of the downloadable filter data.

In order to advantageously offload the burden of unwanted e-mail messages from the bulk of the mail systems of an organisation it is preferred that the filtering mechanisms are in place upstream of the primary mail server.

Viewed from another aspect the invention also provides a computer program product comprising a computer program operable to control a computer to process received e-mail messages, said computer program comprising:

- (i) e-mail filtering logic operable to receive an e-mail message and to apply at least one test to identify a received e-mail message as a potentially unwanted e-mail message; and
- (ii) message forwarding logic operable to forward said potentially unwanted e-mail message to its addressee together with a prompt for said addressee to provide feedback as to whether or not said received e-mail message is an unwanted e-mail message.

The user feedback mechanism applied to suspect e-mail messages is potentially advantageous in its own right independently of the central downloadable source of filter data.

Viewed from a further aspect the invention also provides a computer program product comprising a computer program operable to control a computer to provide downloadable filter data for identifying unwanted e-mail messages, said computer program comprising:

- (i) user report receiving logic operable to receive a user report of an unwanted e-mail message received by said user of said downloadable filter data; and
- (ii) filter data updating logic operable in response to receipt of one or more of said user reports to modify said downloadable filter data to add a test to identify a new unwanted e-mail message.

It will be appreciated that the source of the downloadable filter data itself represents a complementary aspect of the present invention. The downloadable data source and the client system using that downloadable data may be physically separated by considerable distance and may be provided in different countries. Both the client and the data source are separate aspects of the same inventive concept.

Further aspects of the invention provide a method of processing received e-mail messages and an apparatus for processing received e-mail messages.

The above, and other objects, features and advantages of this invention will be apparent from the following detailed description of illustrative embodiments which is to be read in connection with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 schematically illustrates an arrangement of a filter data provider and filter data users;

Figure 2 is a flow diagram illustrating the operation of a client user of the filter data;

Figure 3 schematically illustrates the encapsulation of a suspect e-mail message within a markup language document with voting buttons;

Figure 4 is a flow diagram illustrating the processing by a user of the message of Figure 3;

Figure 5 is a flow diagram illustrating the response of a system to votes received from recipients of the message of Figure 3;

Figure 6 is a flow diagram illustrating the processing applied by the downloadable filtered data provider on receipt of user reports of problem e-mails; and

Figure 7 is a schematic diagram showing a computer that may be used to implement the above described techniques.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 illustrates a plurality of e-mail users in the form of client computers connected via respective mail servers and gateways through the internet. A central provider 1 of the downloadable filter data uses an attached rule database 2 to generate filter data files that may be downloaded by individual connected gateway computers 4, 6, 8. These gateway computers apply the downloaded filter data specified tests to received e-mail messages prior to passing these along to their associated mail servers. The individual gateways 4, 6, 8 may also apply locally defined filtering rules specific to that particular organisation or user.

The rules specified in the rule database 2 may be of an algorithmic form rather than a less flexible direct regular expression form. A predetermined list of words identified as common within unwanted e-mail messages may be established. Words such as “buy”, “free”, “credit” and others have a relatively higher likelihood of occurrence within unwanted e-mail messages than in wanted e-mail messages. By being responsive both to the occurrence of such predetermined rules and the size of the e-mail message itself an indication of the likelihood of an e-mail message being an unwanted e-mail message can be algorithmically determined. Individual words may be given a greater or lesser weighting in the degree to which they indicate that an e-mail message is an unwanted e-mail message. When an e-mail message has been processed by this “fuzzy logic” type testing mechanism, then an indication will be given as to whether or not the e-mail message is definitely unwanted, potentially unwanted or wanted. Wanted e-mail

messages can be passed through to the addressee, definitely unwanted e-mail messages can be deleted and potentially unwanted e-mail messages can be subject to further processing as described below.

The rules may also be responsive to the addressee of a received e-mail message. If a particular e-mail message is detected as being addressed to users who do not exist as well as some that do exist, then this may be indicative of an unwanted e-mail message generated using lexicon based techniques. Depending upon the preponderance of invalid e-mail addresses compared to valid e-mail addresses, such a rule could be used to classify an e-mail message as either definitely unwanted, suspect or wanted.

Figure 2 is a flow diagram schematically illustrating the rule based processing that may be applied by the various gateway computers 4, 6, 8. At step 10, an e-mail message is received. At step 12 the e-mail message is compared with the downloaded and locally generated rule sets held by that gateway computer 4, 6, 8 and scored as to its likelihood of being an unwanted e-mail message.

At step 14, a determination is made from the score as to whether or not the e-mail message is definitely unwanted. If the e-mail message falls within this category, then it is deleted at step 16. If the e-mail message is not definitely unwanted, then it passes to step 18 where a test is made as to whether or not its score indicates that it is a potentially unwanted suspect e-mail message. If the e-mail message is a potentially unwanted E-mail message, then it is passed to step 20 where it is encapsulated within an HTML mail message with voting buttons added to the bottom of the mail message to enable a recipient to provide feedback to a central source as to whether or not that encapsulated mail message is in fact an unwanted mail message. Button is a term that indicates a mechanism within the message allowing automated feedback rather than a specific appearance or coding form.

If the e-mail message is definitely wanted or after encapsulation at step 20, then the message is forwarded to the addressee at step 22.

Figure 3 schematically illustrates a markup language document 24 containing the encapsulated suspect e-mail message 26. The voting buttons 28 provided at the foot of

the message 24 sent to the user allows the user to provide feedback to a central source effectively voting upon the nature of the encapsulated e-mail message 26. Within an individual gateway computer 4, 6, 8, a threshold of a predetermined number of votes positively identifying an e-mail as an unwanted e-mail may be set before triggering a report to the central filter data provider or the generation of a new local rule. The feedback mechanism illustrated is shown in the form of classic HTML buttons, but it will be appreciated that different user interface mechanisms may be provided in conjunction with the encapsulated message to allow a user to provide their feedback as to the nature of the encapsulated E-mail message 26.

Figure 4 is a flow diagram illustrating the processing performed by the recipient of a message such as illustrated in Figure 3. At step 30 the user receives the message. At step 32 the user votes on the nature of the message by clicking on one of the buttons 28. At step 34 this vote is returned to the gateway computer 4, 6, 8 associated with that user.

Figure 5 is a flow diagram illustrating how the gateway computer 4, 6, 8 may respond to votes upon suspect e-mail messages. At step 36 the system waits for votes to be received. When a vote is received, step 38 determines whether or not this newly received vote has the result of making the total number of votes received in relation to that particular encapsulated message 26 exceed a predetermined threshold level, such as three votes positively identifying the encapsulated message 26 as unwanted. If the threshold has not yet been exceeded, then step 40 serves to increment the current count and processing terminates. Processing to accommodate conflicting votes may also be provided.

If the threshold has now been exceeded, then step 42 issues a notification to an administrator of the gateway computer 4, 6, 8. The notification to the administrator generated at step 42 can give an indication of the unwanted e-mail message and allow the administrator to either confirm or not confirm the appropriateness of now treating that e-mail message as unwanted and generating an associated new rule. The administrator makes this confirmation at step 44.

If the administrator indicates that the message should not be treated as unwanted, then step 46 stops further counting of votes relating to that message. If the e-mail message is confirmed as unwanted, then step 48 automatically generates a new local rule to filter out that e-mail message and step 50 provides a notification of the nature of that e-mail message to the central downloadable filter data source such that other users may benefit from the experience of the current user.

It will be appreciated that the confirmation steps by the administrator could be removed and the process once the votes had exceeded the predetermined threshold level could be completely automated. This accords well with the realisation that the determination of whether or not an e-mail message is a Spam e-mail message is one properly decided by the recipients and intervention by an administrator may not be necessary or appropriate.

Figure 6 is a flow diagram illustrating how the central source of downloadable filter data may respond to notifications from separate gateway computers 4, 6, 8 of newly detected unwanted e-mail messages. At step 52, the system waits for new notifications. At step 54, the system checks as to whether or not a newly received notification means that a threshold level of notifications relating to a particular e-mail message has now been received. If the threshold level has not yet been exceeded, then step 56 increments the current count and processing terminates.

If the threshold has been exceeded, then a central authority confirming new globally applicable rules is notified at step 58. Given that new rules set up within the downloadable filtered data will impact potentially all the users of the system, there is a high level of justification for at least having some level of manual checking of new global rules. It may be that the new rules are automatically added to the set and checked retrospectively in order to provide the optimum speed of response. It could be that the confirmation would not be required if severally highly trusted users reported an e-mail message as unwanted compared with perhaps individual users.

If confirmation is being sought, then this is received at step 60. If the new rule is not confirmed, then step 62 terminates further counting in relation to that e-mail

message. If the new rule is confirmed, then step 64 automatically adds it to the downloadable rule set 2.

Figure 7 schematically illustrates a computer 200 of a type that may be used to execute the computer programs described above. The computer 200 includes a central processing unit 202, a random access memory 204, a read-only memory 206, a hard disk drive 208, a display driver 210 and display 212, a user input/output circuit 214, a keyboard 216, a mouse 218 and a network interface circuit 220, all coupled via a common bus 222. In operation, the central processing unit 202 executes computer programs using the random access memory 204 as its working memory. The computer programs may be stored within the read-only memory 206, the hard disk drive 208 or retrieved via the network interface circuit 220 from a remote source. The computer 200 displays the results of its processing activity to the user via the display driver 210 and the display 212. The computer 200 receives control inputs from the user via the user input/output circuit 214, the keyboard 216 and the mouse 218.

The computer program product described above may take the form of a computer program stored within the computer system 200 on the hard disk drive 208, within the random access memory 204, within the read-only memory 206, or downloaded via the network interface circuit 220. The computer program product may also take the form of a recording medium such as a compact disk or floppy disk drive that may be used for distribution purposes. When operating under control of the above described computer program product, the various components of the computer 200 serve to provide the appropriate circuits and logic for carrying out the above described functions and acts. It will be appreciated that the computer 200 illustrated in Figure 7 is merely one example of a type of computer that may execute the computer program product, method and provide the apparatus described above.

Although illustrative embodiments of the invention have been described in detail herein with reference to the accompanying drawings, it is to be understood that the invention is not limited to those precise embodiments, and that various changes and modifications can be effected therein by one skilled in the art without departing from the scope and spirit of the invention as defined by the appended claims.